

NOW

nowdigital.com.au

WHITEPAPER | MARCH 2026

The generative visibility shift

The generative visibility shift

Executive summary

Search has been absorbed. Digital discovery is changing more than it has since PageRank was introduced in 1998. Instead of earning visibility through rankings and clicks, organisations now gain visibility by being included and cited in AI-generated answers.

AI systems no longer just list links. They now generate conclusions.

For Australian organisations, this means a major change in how visibility, authority, and market influence are distributed.

The three structural shifts

1. From ranking to inclusion

Visibility is now about whether your organisation is mentioned in the answer, not where it appears on a results page.

2. From traffic volume to answer authority

Clicks are dropping because AI systems answer questions directly. Now, the real competition is about how often your organisation is included in those answers, not how many clicks you get.

3. From human-only UX to proxy UX

AI agents are now doing more for users, such as extracting content, verifying facts, and completing tasks. If your systems aren't easy for machines to read, your organisation won't be seen by these new users.

The business risk of inaction

Organisations that do not adapt face three main risks:

- **Narrative Displacement** – Competitors become the cited authority in generative responses.
- **Revenue Invisibility** – Your brand fails to appear in categorical discovery (“best provider for...” queries).
- **Trust Erosion** – AI systems synthesise shallow or third-party descriptions of your organisation instead of authoritative source material.

Being a leader in your category is no longer about being on the first page of search results. It now depends on being included in the answer itself.

The new performance hierarchy

Traditional digital metrics are becoming less important. Boards should start tracking how often their organisation is included in AI-generated answers.

Traditional KPI	Generative KPI
Ranking Position	Inclusion Rate in AI Answers
Click-Through Rate	Citation Frequency
Traffic Volume	Engagement Depth from AI Referrals
Bounce Rate	Proxy Task Completion Capability

Your organisation’s advantage depends on how it is set up for machines to understand.

Why this matters in Australia

Australian consumers are using generative AI faster than most of the world. However, many local businesses, especially small and medium ones, are not investing enough in AI infrastructure.

This creates a widening discovery gap:

- Consumers are asking AI for recommendations.
- Australian organisations are not set up to be included in AI-generated answers.
- If local organisations do not change their approach, global competitors will shape the story in the Australian market.

What organisations must do now

To stay competitive in this new AI-driven discovery space, organisations should:

1. Architect structured authority through entity-first content design.
2. Optimise for atomic clarity and semantic retrievability.
3. Ensure infrastructure supports server-side rendering and machine accessibility.
4. Align traditional SEO authority with AEO/GEO machine inclusion.
5. Redesign post-click environments for AI proxy navigation.

This requires a full redesign of your infrastructure and authority.

The strategic imperative

This shift to generative AI is a fundamental change, not just a test or experiment.

Organisations that act early will control structured authority in their category.

Organisations that wait too long will only get leftover visibility.

NOW Digital partners with Australian organisations ready to move from search participation to generative dominance.

The next step is to review how your organisation is included, how your infrastructure operates, and how you are positioned to exercise authority in the new answer-driven economy.

This change is already underway. The real question is whether your organisation is ready for it.

Research Whitepaper

Optimising content for generative answer machines

The digital search landscape is undergoing one of the most transformative changes in decades: the shift from Search Engine Optimisation (SEO) to Answer Engine Optimisation (AEO) and Generative Engine Optimisation (GEO).

Since Google's PageRank in 1998, digital content strategy has operated in an index economy. Optimisation has focused on metrics like searches, clicks, impressions, and consumption to assess page effectiveness.

With the rapid adoption of LLM (Large Language Model) and LMM (Large Multimodal Model) technology, **AI agents are increasingly assuming user identities**. Optimisation shifts from indexing and retrieval to a synthesis economy.

- Traditional metrics (clicks and rankings) are losing significance, as executives must now prioritise synthesis, measured by the quality of answers and citations.
- AEO/GEO factors now overshadow SEO in strategic planning.
- Executive value now lies in the quality of AI-generated outputs, not search engine rankings.

AI agents as users

As AI agents become primary users, browsing shifts from optimising user experience to optimising data retrieval and delivery. This change is immediate and measurable, with research showing:

- The estimated traditional search volume for this year will decline by up to 25% within 24 months¹.
- Instances of traffic referred via LLM and LMM processing are experiencing significant jumps, with some instances increasing from 17,076 to 107,100 (527%) from January to May 2025 alone².

The criterion bolstering this operational strategy shift is the movement away from CTR (Click-Through Rate) and towards SOA (Share-of-answer) as a key metric. Here, brand utility becomes defined not by ranking, but by citation and synthesis.

Economy of index to economy of synthesis

Traditional SEO strategy relies on a contract that establishes that search engines crawl brand-provided information to index content and organise data. In exchange for this information, search engines distribute traffic to their publishers.

Generative AI disrupts the traditional search contract. Previously, the AEO/GEO strategy page or brand relevance was determined through SEO processes, including crawling and indexing to construct an inverted index. From there, query matching proceeded, based on lexical relevance and a range of authority and trust signals³.

Within AEO/GEO strategy, AI tools go beyond information indexing. Search inputs routed through LLMs/LMMs are used to generate novel outcomes, which are distinguished by their interaction at the semantic level, prioritising qualifiers like meaning and intent.

This surpasses traditional, or even modern 'AI' search engine interaction with Machine Learning, which reliably peaks at 'learn-to-rank' models that still fundamentally offer only ranked document lists.

The generative difference occurs in these distinctions:

- Traffic is no longer promised in exchange for content data.
- Outputs are synthesised rather than retrieved and listed.
- Intent takes priority over lexical comparison.

Zero-click economy presents itself as a symptom of this structural shift, acting both as a circuit breaker between the approaches to content optimisation and the phenomenon that allows us to justify a commitment to the transition.

Zero-click refers to a user search behaviour in which a user's information needs are met without requiring any click-through after the search. No external link is necessary to fulfil the query, as the query is adequately addressed within the search results themselves.

This change is measurable. Industry leaders forecast that traditional search volume will fall 25% by late 2025, while AI chatbots and agents will rise. When Google SGE or 'AI Overview' appears, organic CTR drops up to 43.5% on desktop and 34.5% on mobile, even in 'best-case scenarios'⁴.

¹ <https://www.gartner.com/en/newsroom/press-releases/2024-02-19-gartner-predicts-search-engine-volume-will-drop-25-percent-by-2026-due-to-ai-chatbots-and-other-virtual-agents>

² <https://searchengineland.com/ai-traffic-up-seo-rewritten-459954>

³ https://assets.cambridge.org/97805218/65715/frontmatter/9780521865715_frontmatter.pdf

⁴ https://assets.cambridge.org/97805218/65715/frontmatter/9780521865715_frontmatter.pdf

This is a system reallocation, not search engine obsolescence. Users still search to navigate, but informational functions are absorbed into the interface.

This rapid drop in CTR implies that the user is, too, assuming a new identity. The act of browsing is taken off the user, and verification takes its place. 'SERP-mediated investigative behaviour' starts facing declines, because the verification process is being handled by the system rather than placing a burden on the user.

Although overall referral volumes from AI assistants are lower, engagement from these sources has shown to be significantly higher. This highlights the necessity of executives to track engagement quality, not just volume, when evaluating digital performance.

Referrals from AI assistants, such as Perplexity.ai, have an average referral page engagement of 552 seconds (~9 minutes).

When compared with the average engagement time on a page referred from Google.com (~1.5 minutes)⁵, it is evident that competing for site engagement necessitates a shift from optimising for value volume to optimising for value intention.

Mechanics of machine knowledge

To optimise for ML, distinguish human cognition from computation. AI assistants don't read text; they calculate it. AI transforms language into coordinates in high-dimensional vector spaces, measuring the distance between meanings rather than reading sentences.

In a similar way to the necessity of a grasp of SEO architecture for optimising content for search engine relevance in the context of information crawling and indexing, a grasp of agentic extraction is necessary for optimising content for AI competency.

This involves understanding agentic extraction as a transition from 'Link-Graph' to 'Meaning-Graph' information processing.

This is where content relevance is no longer determined by inferred link strength between documents. In AI retrieval contexts, the web is represented as a graph of semantic embeddings, connected by meaning strength.

AI crawlers may miss vital information in content optimised for SEO and hidden by CSR. This 'detection gap' occurs when content fails to match AI models' relevance criteria.

- Content is assessed on vector proximity.
- Extraction replaces indexing.
- Content that relies on legacy rendering risks becoming invisible within a meaning-graph economy.

RAG architecture

When enterprise ML Search Engines (such as Copilot, Perplexity or SGE) incorporate AI assistants, a framework named 'RAG (Retrieval-Augmented Generation) Architecture' is utilised to compensate for 'knowledge cut-off'.

Because Large language models are trained on static datasets, without inherent, external retrieval; they operate within frozen knowledge boundaries.

⁵ <https://seranking.com/blog/ai-traffic-research-study/>

RAG Architecture connects these static models to live data. These services like SGE (Google AI Overview) retrieve relevant web content in real time and inject it into a ML/MLL generation process. This introduces a dynamic facet, which opens the knowledge cut-off and allows for outputs grounded in current, verifiable content sources.

RAG determines what AI systems read, select and synthesise from your online content.

The RAG Architecture process is comprised of four distinct processes that provide valuable insights into exactly how online content is treated by AI assistants, and how content can be optimised for selection.

1. Vectorisation

Unstructured data, such as common text, or page documents undergo ingestion and semantic index. This prepares the content for the system.

Technical Action

documents are divided into smaller chunks, representing 200-500 “tokens”. This aims to preserve the holistic meanings of separate expressions by attempting to avoid diluting its understanding through the inclusion of unrelated context.

Unlike traditional indexing systems, which catalogue and organise text based on keywords, embedding models convert that text into high-dimensional vectors that numerically represent semantic relationships. Here, the relative semantic connection between words is understood in an array of numerical expressions⁶.

Latent mathematical space is used to express the closeness of meaning, where concepts with more shared meaning are geometrically closer⁷.

Strategic Takeaway

Content must be modular and thematically focused: Geometric proximity now outweighs keyword density.

- One clear idea per block
- Avoid conceptual stacking
- Maintain granular structure

2. Retrieval

Once a query is submitted, the system does not merely attempt to associate or pair relevant keywords. Instead, it performs a Mathematical Similarity Search (MSS) with the re-represented vector data.

Technical Action

The submitted query is subjected to the same embedding models to verify it exists within matching coordinate systems⁸ (deem it relevant). An Approximate Nearest Neighbour (ANN) search is used, which calculates the mathematical distance (such as Euclidean distance or cosine similarity) between the vector of the query and that of the stored

⁶ <https://arxiv.org/pdf/2501.14342>

⁷ <https://redis.io/blog/vector-similarity/>

⁸ <https://arxiv.org/pdf/2501.14342>

documents⁹. So, even if they do not share exact keywords, the system is retrieving the “top-k” expressions that express the most semantic relation to the query¹⁰.

Strategic Takeaway

Vague content produces weak vectors. Precise content produces retrievable authority.

- Use semantic variations and natural phrasing
- Anticipate question-based language (“Why should I..”)
- Saturate content with concrete entities (names, dates, locations)
- Keep problem and solution within the same semantic window

3. Augmentation

This phase allows the retrieval and generation stage to communicate by preparing the input for the LLM through context integration.

Technical Action

Having retrieved the most relevant data from the external source, the users original query is re-integrated into a predefined prompt template, and an enriched prompt is formed¹¹. The system creates a structured prompt, such as,

“Context: [Retrieved Data]. Question: [User Query]. Answer:”

This explicitly instructs the LLM to utilise the enriched information to generate a highly relevant prompt¹².

Strategic Takeaway

Optimise for context-window efficiency. If a section relies on surrounding narrative, it risks losing coherence during injection

- Use clear syntax (Subject-Verb-Object)
- Avoid layout-dependent meaning
- Ensure each block stands alone
- Increase information density

4. Generation

The model generates a response using both its pre-trained knowledge and the retrieved content.

Technical Action

by conditioning the generation with retrieved content, the system displays non-parametric memory by grounding its answer in specific, semantically verifiable found information rather than being restricted to a reliance on its pre-trained, or parametric data¹³. Vectors replace keywords, Answers replace links, and content demands for ranking is replaced by demands for clarity, authority, and structure.

⁹ <https://arxiv.org/html/2507.18910v1>

¹⁰ <https://teambrookvale.com.au/articles/rag-technical-deep-dive>

¹¹ <https://www.ibm.com/think/topics/retrieval-augmented-generation>

¹² Ibid.

¹³ <https://arxiv.org/html/2507.18910v1>

Strategic Takeaway

Write for synthesis.

- Use definitive language
- Avoid hedging where certainty exists
- Structure internally as Problem → Solution → Proof
- Signal authority consistently

Definitive content produces definitive answers.

Generative engine optimisation statistics insight

A collaborative research team from Princeton University, Georgia Tech and Google DeepMind released a study aimed at discovering which optimisation strategies forced an AI to cite a specific site. Within this process, GEO-bench was introduced: a framework designed to rigorously categorise the content characteristics that drive visibility across tens of thousands of generative responses.

The following section identifies and contextualises some of the most valuable insights¹⁴

Citation of sources

Integrating authoritative data, such as government, academic, or industry-leader sources, signals a higher truth-probability

- **30-40% relative improvement** on Position-Adjusted Word Count
- Significant visibility boosts for lower-ranked pages, with minimal effort

Addition of quotations

Direct quotes add necessary authenticity and depth, particularly for domains involving society or culture.

- **15-30% improvement** on Subjective Impression metrics
- Elevates perceived quality and contributes to this “stickiness” in generation

Fluency Optimisation

Intentionally simplifying language and removing semantic ambiguities reduces the friction of readability during retrieval.

- **15-30% overall boost** in page visibility
- Content parity is now considered a content quality marker.

The keyword stuffing penalty

Keyword stuffing and other similar legacy content optimisation are observed to communicate a low-quality, or shortcut foundation to LLMs/LMMs

- **10% worse performance** than completely SEO optimised content

Domain-specific efficacy

¹⁴ <https://arxiv.org/html/2311.09735v3>

The statistics reveal interesting industry-contextual nuances. The efficacy of these strategies demonstrated significant variability across industries and content styles.

Method	Relative Improvement (%) in Visibility					Method	Top Performing Tags		
	Rank-1	Rank-2	Rank-3	Rank-4	Rank-5		Rank-1	Rank-2	Rank-3
Authoritative	-6.0	4.1	-0.6	12.6	6.1	Authoritative	Debate	History	Science
Fluency Opt.	-2.0	5.2	3.6	-4.4	2.2	Fluency Opt.	Business	Science	Health
Cite Sources	-30.3	2.5	20.4	15.5	115.1	Cite Sources	Statement	Facts	Law & Gov.
Quotation Addition	-22.9	-7.0	3.5	25.1	99.7	Quotation Addition	People & Society	Explanation	History
Statistics Addition	-20.6	-3.9	8.1	10.0	97.9	Statistics Addition	Law & Gov.	Debate	Opinion

Table 2. Visibility changes through GEO methods for sources with different Rankings in Search Engine. GEO is especially helpful for lower ranked websites.

Table 3. Top Performing categories for each of the GEO methods. Website-owners can choose relevant GEO strategy based on their target domain.

<https://arxiv.org/html/2311.09735v3>

Strategy	Top-Performing Domain	Visibility Gain (Avg)	Mechanism of Action
Statistics Addition	Law and Government	+41%	Concrete data reduces model entropy.
Source Citation	Science & Factual	+40%	Secondary authority signals verification.
Quotation Addition	History & Society	+38%	Expert quotes add depth and authenticity.
Fluency Optimisation	Across All Domains	+29%	Reduces computational friction during retrieval.

Australian AI adoption insights and impacts

The SME productivity gap

Australia's economy is hindered by a two-speed reality in which a lagging SME sector accounts for almost 70% of private employment¹⁵. This fails to match the digitisation of large-scale enterprises.

- Australian SMEs are running at about **50% less productive** than their large firms¹⁶. This creates a 'productivity gap' that disqualifies us from sitting with "Advanced Economies" such as the UK (16% productivity gap), Israel (31%), Portugal (34%) and Germany (39%)¹⁷.

¹⁵ <https://doi.org/10.1787/fa521246-en>

¹⁶ <https://publicaccountant.com.au/features/shrinking-sme-productivity-gap>

¹⁷ <https://www.mckinsey.com/featured-insights/week-in-charts/small-enterprises-big-productivity-gaps>

- Poor "diffusion of innovation" prevents the broader economy from capturing the projected **\$112B annual AI opportunity** by 2030¹⁸.

The investment deficit

While adoption rates appear superficially similar to those of global leaders, Australia suffers from chronic underinvestment in the underlying infrastructure required to turn AI into a functioning productivity engine.

- In recent years, the United States has experienced productivity rate growth of 1.7-2.5% p.a, compared the Australia's 0.7%¹⁹
- Although Australia and the US have similar adoption rates for corporate AI. The driver between the productivity difference is clear when we see that the US 10-11x more on relative investment in AI enablers and ground-up infrastructure²⁰

The ROI disconnect

There is a stark misalignment between enterprise spending and actual results, as most local organisations prioritise internal efficiency over external revenue growth.

- Despite budgets reaching **\$28M, 72% of CDAOs** report that current AI strategies have failed to meet ROI expectations²¹.
- A strategic fixation on efficiency (57%) over revenue growth (25%) is stalling the transition to an agentic economy²².

Australian consumers pull as Australian businesses push

Australian consumers are adopting GenAI faster than their global peers, yet local businesses, in particular, SMEs, are failing to appear in the discovery layer.

- **49% of Australians** use GenAI (surpassing the US/UK), and **71%** have replaced traditional search with AI for retail recommendations²³.
- Consumers are asking the questions, but with only **19% of businesses** investing in AI infrastructure²⁴, local brands are not being included in the answers.

Barriers to integration

Adoption is being throttled by a cautious "wait and see" approach, seemingly driven by imminent financial and risk-management considerations

- The primary drivers of hesitancy are **Security (29%)** and **Budget (28%)** concerns²⁵.
- This defensive posture creates a market vacuum that global, infrastructure-prepared competitors are demonstratively, quickly filling.

¹⁸ <https://aiia.com.au/wp-content/uploads/2025/10/Australias-AI-Opportunity-Report.pdf>

¹⁹ <https://www.atse.org.au/media/vkll0utx/251210-unleashing-growth-australias-ai-investment-blueprint.pdf>

²⁰ Ibid.

²¹ <https://research.adapt.com.au/wp-content/uploads/2025/08/State-of-the-Nation-2025-Data-and-AI-in-Australia.pdf>

²² <https://www.decidr.ai/national-ai-readiness-index-report-2025>

²³ <https://www.undivided.com.au/blog/optimising-for-generative-ai-the-future-of-digital-marketing-in-australia>

²⁴ <https://www.decidr.ai/national-ai-readiness-index-report-2025>

²⁵ Ibid.

Entity First Strategic Framework

Authored identity vs. 'Sensical nonsense'

The problem

When content strategy relies on a traditional SEO-oriented approach, it creates a '**detection gap**' for AI assistants. If site content is optimised for metadata and clickstream, rather than for AI synthesis; users may experience AI citing information associated with a brand presenting grammatically correct, but factually shallow - sensible nonsense that lacks true brand authority.

The strategy

Implement a modern Single Source of Truth (SSoT). By moving beyond metadata to approaches such as latent schema marking can be used as part of a GEO-facing SSoT attempts to ensure content data is centralised, authoritative, accessible and integrated, allowing to co-opt broader global authority by allowing AI to pull from pre-verified and structured descriptions, rather than third-party considerations.

Content decomposition and atomic facts

AI systems do not process data linearly, as a human reading text, but instead process via 'fact decomposition'. This process, whereby text is broken down into verifiable 'claim' units, allows the system to establish 'atomic facts'²⁶.

To ensure consistent outcomes, it is best to verify that facts are grounded, before they even reach the AI. In other words, paying special consideration to indexed text, to limit the opportunity of failure in the retrieval or synthesis of those facts.

Content needs to contain a high subclaim atomicity, meaning it's written in such a way that it reinforces consistent meaning or truth throughout itself, rather than relying on a holistic or totalising interpretation to fully draw out its semantic content²⁷.

The principle: AI processes data via 'fact decomposition', whereby text is broken into verifiable "claim units" or Atomic Facts.

The implementation: Use Entity-Description Pairs (EDPs) to express optimised retrieval units (e.g., "NOW Digital Pricing" + specific factual descriptor).

The impact: This 'context-efficiency' enables systems to traverse huge datasets, without context overwhelm. These extracts claim correctness prior to the synthesis stage.

Operationalising atomicity

- **High Subclaim Atomicity:** Content must be written to reinforce consistent meaning throughout, rather than relying on a holistic interpretation.

²⁶ <https://arxiv.org/html/2503.19574v1>

²⁷ <https://arxiv.org/pdf/2503.15354?>

- **The Risk:** Paragraphs with low atomicity, dense language, or convoluted clauses confuse the system’s truth-verification criteria, leading to lowered confidence scores and potential hallucinations.

Inverted pyramid processing

- **Position Bias:** Research from Stanford and UC Berkeley confirms the ‘lost-in-the-middle’ phenomenon, where AI retrieval favours information at the direct start or end of a section²⁸.
- **Answer-First Architecture:** Place at least one citable atomic fact within the first 50-60 words.

Example: Replace narrative fluff (“We are committed to...”) with pragmatic, actionable units (“To renew, account holders must...”). This mimics the attention-weight allocation of machine encoding²⁹.

Infrastructure: SSR vs. CSR

- **The technical gap:** AI crawlers like GPTBot and Claudebot struggle with JavaScript. Sites reliant on Client-Side Rendering (CSR) face a detection gap of 27-29%³⁰.
- **The Requirement:** Deliver content via Server-Side Rendering (SSR) or Static Site Generation (SSG). A full DOM and structured data must be visible in the initial HTML response to ensure entity detection.

Latent schema and @Graph architecture

- **Narrative mapping:** Use ‘@graph’ architecture to compress nodes (Article, Author, Organisation) into a coherent narrative.
- **SameAs disambiguation:** Use the SameAs markup to link internal entities to external knowledge bases such as Wikidata or Crunchbase, attaching a global identifier that co-opts authority and prevents ambiguity.
- **The outcome:** Confident, unambiguous entities, (for example, explicitly distinguishing **NOW Digital** from unrelated entities such as **9Now** or the **National Organisation for Women**) thereby increasing the confidence score of the generative response³¹.

The failure of explicit workarounds

- **GTM Injections:** Injecting schema via Google Tag Manager is increasingly ineffective. Generative engines prioritise implicit semantic processing and often crawl before JavaScript renders.
- **Semantic Gaps:** When schema claims contradict contextual signifiers (like user reviews), the engine identifies a semantic gap and deprioritises the source.

²⁸ (<https://arxiv.org/pdf/2307.03172>)

²⁹ [10.18653/v1/2024.acl-short.35](https://arxiv.org/pdf/10.18653/v1/2024.acl-short.35)

³⁰ [10.63282/3050-9416.IJAIBDCMS-V7I1P112](https://arxiv.org/pdf/10.63282/3050-9416.IJAIBDCMS-V7I1P112)

³¹ <https://eugeneasahara.com/2025/10/10/stories-the-unit-of-human-level-intelligence>

- **Agentic Failures:** If an actionable task (e.g., a "Buy" button) is not semantically labelled in the HTML, the AI proxy cannot "see" the button and act on the user's behalf.

Bifurcated approach

The new metric hierarchy

The shift toward AEO/GEO requires a two-pronged optimisation strategy that balances human-centric SEO with machine-readability. In this environment, ranked-first outcomes are secondary to inclusion and citation rates.

- Inclusion in AI-generated answers is now a primary KPI; visibility is no longer defined by a click-through, but by status as a cited anchor of certainty.
- Brands must transition from being seekers of search traffic to becoming the authoritative sources that AI systems rely upon to validate their own outputs.

Prime signalling and E-E-A-T

To secure citation in answer engines, content must be structured using Prime Signalling. These are intentional technical and narrative markers that encourage machine trust.

Experience	<ul style="list-style-type: none"> • First-hand content like case studies or specific projects. • Demonstrates that the content is not a generic synthesis but situated.
Expertise	<ul style="list-style-type: none"> • Consistent use of sector terminology • Established methodology (consistent with field standards) • Explicit logical approach
Authoritativeness	<ul style="list-style-type: none"> • Strong citations (peer-reviewed, published by industry leader, etc.) • External, corroborative style of writing, signalling relational authority, rather than rhetorical
Trustworthiness	<ul style="list-style-type: none"> • Explicit scope and limitation consideration • In-text references (inline sourcing) • Timestamps, transcripts

Leveraging the **E-E-A-T framework** provides the structural integrity necessary for AI models to verify and cite content.

- Content that fails to demonstrate these four pillars is effectively invisible to generative engines, regardless of its traditional keyword relevance.

Community validation

Many prevalent generative engines display a strong preference for community-moderated platforms, using them as proxies for real-world trust and brand validity.

- Platforms like **Wikipedia and Reddit** account for **25% to 40% of generative answers** across various AI models³².

³² <https://pressonify.ai/blog/ai-search-platform-comparison-2026>

- It is now imperative that brands harness the benefits of community-led sites in developing brand narratives within the generative sphere. These platforms are the high-trust signals that dictate how AI models perceive and present your organisation.

Bridging the discovery gap

There is a massive disconnect between how AI recognises explicit brands versus how it handles categorical queries.

- While LLMs show **99.4% precision** for explicit brand names, recognition drops to as low as **3.32% for categorical queries**³³ (e.g., "What is the best enterprise product for X?")
- This gap reveals the necessity of the bifurcated approach. Traditional SEO trust signals establish an information footprint, while AEO/GEO optimisation translates that footprint into machine-digestible material that triggers citations.

The SEO + AEO/GEO nexus

Optimising in a vacuum is a recipe for failure. Authentic, foundational visibility (SEO) is the prerequisite for establishing the data points an LLM needs to cross-validate a brand.

- **The Goal:** Use traditional SEO to build the authority, and AEO/GEO to ensure that authority is correctly interpreted, cited, and recommended in the generative discovery layer.
- **Impact:** By aligning human-oriented trust signals with machine-oriented technical structures, businesses ensure they are not just "findable," but are the preferred answer in the agentic economy.

Post-click UX: proxy design for AI

The death of digital brochure

Corporate and government web design have historically been guided by SEO logic to create "digital brochures": interfaces that prioritise visual cleanliness and narrative-driven layouts to capture human attention.

The shift

Design choices optimised for human psychology (e.g., whitespace, brand storytelling) are often opaque to machine attention. As we move into an agentic economy, the "User" is no longer exclusively human.

Agentic proxies

Task-specific AI agents are increasingly integrated within enterprise applications to automate workflows. By acting as a **proxy user**, agents extract content and execute tasks in the post-click environment.

³³ <https://arxiv.org/pdf/2601.00912>

- Adoption rates for these AI-based, task-specific entities are expected to grow from **5% of enterprise applications** featuring these agents in late 2025 to **40% by the end of 2026**³⁴.

Optimising for machine perception

The challenge: Adapting strong considerations towards the non-perceptual aspects of content consumption (extraction) for non-human entities³⁵.

If an AI agent is unable to “read” a workflow due to poor semantic labelling or a narrative-heavy layout, the automated process fails.

The strategic pivot

A transition from designing for purely human eyes to adopting a Proxy UX strategy. This ensures that when an AI agent visits a site to perform a task (such as processing a renewal or extracting data), it encounters a structured, machine-navigable environment.

Machine Trust

Research into trust and machine systems demonstrate that increased user trust is primarily, correlated with decreased monitoring behaviour³⁶. In other words, users who display high trust in a system are significantly less likely to have the desire to verify its outputs (such as checking citations).

- Monitoring therefore signals scepticism rather than engagement.

To optimise interaction, content must be structured so AI can deliver accurate answers from available information. As trust increases, verification declines and friction reduces.

The following section outlines key considerations in creating high-trust content foundations.

Calibrated trust and the risk of misuse

An optimal use state exists when user reliance aligns with system capability³⁷. Trust erodes when this balance shifts.

- Over-trust leads to misuse, while under-trust results in disuse.

For corporate and government organisations, avoiding both outcomes is critical. Data integrity must align with interface authority to ensure AEO/GEO optimisation does not amplify outdated or inaccurate information.

The trust-vulnerability and knowledge-trust concerns

Higher trust, particularly naive trust, can increase vulnerability. When guardrails are lowered, the attack surface expands and exposure to information risk grows³⁸.

³⁴ <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>

³⁵ <https://link.springer.com/article/10.1007/s11023-018-9482-5>.

³⁶ <https://arxiv.org/pdf/2501.01303>

³⁷ <https://ls9-www.cs.tu-dortmund.de/publications/CHI2023.pdf>

³⁸ <https://arxiv.org/html/2510.18563v1>

Security-first UX must accompany seamless design to ensure confidence is supported by strong governance. Technical familiarity alone does not increase trust.

Users are more confident when outputs visibly link to identifiable data sources, reducing concerns about inaccuracy, without requiring technical explanation³⁹.

The citation-trust inversion

- The presence of citations increases perceived trust yet verifying them decreases it⁴⁰.

Users want citations to exist as reassurance but prefer not to rely on them. This inversion reinforces the shift from SEO to AEO/GEO.

Authority must be embedded within generated answers, with citations visible but secondary to the primary content.

³⁹ <https://arxiv.org/pdf/2501.01303>

⁴⁰ Ibid.